

**Amendments to the Claims:**

This listing of claims will replace all prior versions and listings of claims in the application.

**Listing of Claims:**

1. (Previously Presented) A method for reporting performance information as perceived by a client in a client server system, the method comprising:

detecting that a server is enabled to receive the performance information, wherein the performance information comprises at least one of: a request/response latency, a request/response error code , and a request/response error code frequency; dispatching performance context information from the client to the server, wherein the performance context information comprises one or more of: a client computer system host name, a client user name, a client network adaptor name, a client network adaptor speed, and a client network protocol address;

dispatching a first request from the client to the server, wherein the first request specifies a first remote procedure call (RPC), and wherein the first RPC comprises at least one argument for executing a program on the server;

in response to the first request dispatched from the client to the server, receiving a first response from the server, wherein the first response comprises a result of the first RPC;

in response to the received first response, measuring a time delay from the client's dispatch of the first request to the client's receipt of the first response from the server, wherein the time delay is a number of milliseconds between sending the first request and receiving the first response, and wherein the time delay is a request/response latency of the first request;

storing the first request/response latency at the client until a second request specifying a second RPC is generated by the client;

appending the first request/response latency to a header of the second request from the client to the server; and

dispatching, from the client to the server, the second request with the header, wherein the second request specifies the second RPC, and wherein the second RPC is different from the first RPC.

2. (Original) The method of claim 1, wherein the client is a messaging client, and wherein the server is a messaging server.

3.-8. Canceled.

9. (Previously Presented) The method of claim 1, wherein the client is an email client, and wherein the server is an e-mail server.

10. (Previously Presented) A software program embodied on a computer storage medium, wherein the software program is executable to perform the method of claim 1.

11. Canceled.

12. (Previously Presented) A computer-implemented method, comprising:  
detecting that a server is enabled to receive performance data, wherein the performance data comprises at least one of: a request/response latency, a request/response error code, and a request/response error code frequency;

sending performance context information from a client to the server, wherein the performance context information comprises one or more of: a client computer system host name, a client user name, a client network adaptor name, a client network adaptor speed, and a client network protocol address;

sending a first request from the client to the server, wherein the first request specifies a first remote procedure call (RPC), and wherein the first RPC comprises at least one argument for executing a program on the server;

recording, at the client, a request initiation time for the first request;

receiving, at the client, a first response from the server comprising a result of the first RPC;

in response to the received first response, recording, at the client, a response received time for the first response;

calculating a first request/response latency comprising a difference between the response received time for the first response and the request initiation time for the first request, wherein

the first request/response comprises a number of milliseconds between recording the request initiation time for the first request and recording the response received time for receiving the first response;

storing the first request/response latency at the client until a second request specifying a second RPC is generated by the client; and

sending a second request from the client to the server, the second request comprising:

a header comprising the first request/response latency; and

the second RPC, wherein the second RPC is different from the first RPC, and wherein the second RPC comprises at least one argument for executing a program on the server.

13.-20. Canceled.

21. (Previously Presented) The method according to claim 12, wherein the performance data context information comprises at least one identifier of the performance data context information.

22. Canceled.

23. (Previously Presented) The method according to claim 12, wherein the performance data context information further comprises one or more of the following:

a server computer system host name;

a server network domain name; and

a server type.

24. (Previously Presented) The method according to claim 12, wherein the performance data context information further comprises one or more of the following:  
a globally unique identifier (GUID) associated with a client computer operating system process;  
and  
information regarding the client computer operating system process.

25. (Previously Presented) The method according to claim 24, wherein the performance data context information further comprises a globally unique identifier (GUID) associated with a client server communications session.

26. (Original) The method according to claim 12, wherein the second request further comprises an indication of at least one service desired of the server by the client.

27. (Previously Presented) A computer storage medium having thereon computer executable instructions for performing the method according to claim 12.

28.-29. Canceled.

30. (Previously Presented) A computer-implemented method, comprising:  
detecting that a first server is enabled to receive performance data, wherein the performance data comprises at least one of: a request/response latency, a request/response error code, and a request/response error code frequency;

sending performance context information from a client to the first server, wherein the performance context information comprises one or more of: a client computer system host name, a client user name, a client network adaptor name, a client network adaptor speed, and a client network protocol address;

sending a first request from the client to the first server, wherein the first request specifies a first remote procedure call (RPC), and wherein the first RPC comprises at least one argument for executing a program on the first server;

recording, at the client, a request initiation time for the first request;

receiving, at the client, a first response from the first server comprising a result of the first RPC;

in response to the received first response, recording, at the client, a response received time for the first response;

calculating a request/response latency for the first request/response pair comprising a difference between the response received time for the first response and the request initiation time for the first request;

storing the request/response latency for the first request/response pair at the client until a second request specifying a second RPC is generated by the client; and

    sending a second request from the client to a second server, the second request comprising:

        a header comprising the request/response latency for the first request/response pair; and  
        the second RPC, wherein the second RPC is different from the first RPC, and wherein the second RPC comprises at least one argument for executing a program on the second server.

31. (Previously Presented) The method according to claim 30, further comprising:  
    receiving the second request at the second server;  
    parsing the first request/response latency from the second request; and  
    updating, with the first request/response latency, at least one computer system memory resident performance data accumulator associated with the second server.

32. (Previously Presented) A computer-implemented method, comprising:  
    detecting that a server is enabled to receive performance data, wherein the performance data comprises at least one of: a request/response latency, a request/response error code, and a request/response error code frequency;  
    sending a first request from a client to the server, wherein the first request specifies a first remote procedure call (RPC), and wherein the first RPC comprises at least one argument for executing a program on the server;

    recording, at the client, a request initiation time for the first request;  
    receiving, at the client, a first response from the server comprising a result of the first RPC;

    in response to the received first response, recording, at the client, a response received time for the first response;

    calculating a request/response latency for the first request/response pair comprising a difference between the response received time for the first response and the request initiation time for the first request;

storing, at the client, performance data associated with the first request/response pair and a performance data storage time, the performance data comprising the request/response latency for the first request/response pair;

sending a second request from the client to the server, wherein the second request specifies a second RPC, and wherein the second RPC is different from the first RPC; and

when a difference between a request initiation time for the second request and the performance data storage time associated with the first request/response pair is less than a maximum performance data age threshold, incorporating the performance data associated with the first request/response pair into a header of the second request.

33. (Original) The method according to claim 32, further comprising receiving the maximum performance data age threshold from the server.

34.-41. Canceled.

42. (Previously Presented) A computer system, comprising:

a processor, wherein the processor processes instructions for reporting performance data as perceived by a client in a client server network; and

a computer storage medium having stored thereon the instructions for reporting performance data, the stored instructions comprising a data structure, the data structure comprised of:

a performance data remote procedure call (RPC) extension tag included in an RPC request for prefacing network performance data associated with a client, said performance data RPC extension tag comprising:

a performance data format version field;

at least one performance data format flag;

a performance data size field; and

an uncompressed performance data size field; and

at least one performance data block relating to the performance data RPC extension tag, said at least one performance data block comprising:

a performance data block header, comprising:

a performance data block size field;

a performance data block format version field; and

a performance data block type field; and

a performance data block body comprising the network performance data associated with the client, wherein the performance data RPC extension tag and the at least one performance data block are incorporated in an RPC in each request/response cycle of the transmission of the RPC.

43. (Previously Presented) The computer system of claim 42, wherein the at least one performance data format flag is selected from a group consisting of compress performance data and obfuscate performance data.

44. (Previously Presented) The computer system of claim 42, wherein the performance data block body comprises: a variable byte array storage area; and at least one string type field comprising an offset to a location in the variable byte array storage area where the value of the string type field is stored.

45. (Previously Presented) The computer system of claim 42, wherein the performance data block type field indicates that the performance data block body contains client information performance data context pertaining to the client in the client server network, and wherein the performance data block body comprises:

a client-generated client information performance data context identifier;

a name of a computer system hosting the client;

a name of a user utilizing the client;

a network protocol address associated with the client;

a name of a network adapter associated with the client; and

a speed of the network adapter associated with the client.

46. (Previously Presented) The computer system of claim 42, wherein the performance data block type field indicates that the performance data block body contains server information performance data context pertaining to a server in the client server network, and wherein the performance data block body comprises:

a client-generated server information performance data context identifier;  
a name of the server;  
a network domain name of the server; and  
a server type associated with the server.

47. (Previously Presented) The computer system of claim 42, wherein the performance data block type field indicates that the performance data block body contains client-generated performance data pertaining to a successful request/response pair between the client and a server in the client server network, and wherein the performance data block body comprises:

a client-generated request identifier associated with the successful request/response pair;  
a round trip network latency corresponding to a difference between a time at which the client initiated the successful request/response pair and a time at which the client received the response; and  
a server processing time corresponding to a difference between a time at which the server received the request of the successful request/response pair and a time at which the server initiated the response.

48. (Previously Presented) The computer system of claim 47, wherein the performance data block size is 14 bytes.

49. (Previously Presented) The computer system of claim 42, wherein the performance data block type field indicates that the performance data block body contains client-generated performance data pertaining to a failed request/response pair between the client and a server in the client server network, and wherein the performance data block body comprises:

a client-generated request identifier associated with the failed request/response pair;  
a time to fail corresponding to a difference between a time at which the client initiated the failed request/response pair and a time at which the client determined that the request had failed; and  
a failure code corresponding to a reason for the failure of the request.

50. (Previously Presented) The computer system of claim 42, wherein the performance data block type field indicates that the performance data block body contains server performance data preference information for a server in the client server network, and wherein the performance data block body comprises:

an indication of whether to send client-generated performance data to the server pertaining to client communications with the server;

an indication of whether to send client-generated performance data to the server pertaining to client communications with other servers in the client server network; and

a performance data age threshold beyond which client-stored performance data should not be sent to the server.

51. (Previously Presented) The method of claim 1, wherein after detecting that the server is enabled to receive the performance information, the client receives performance information preferences from the server, and wherein the performance information preferences comprise at least one of: under what circumstances the client should send performance information, whether the client should send performance information related to communications with servers other than the server, and a maximum performance information age threshold over which the client should not send performance information to the server.

52. (Previously Presented) The method of claim 12, wherein after detecting that the server is enabled to receive the performance data, the client receives performance data preference information from the server, and wherein the performance data preference information comprises at least one of: under what circumstances the client should send performance data, whether the client should send performance data related to communications with servers other than the server, and a maximum performance data age threshold over which the client should not send performance data to the server.

53. (Previously Presented) The method of claim 30, wherein after detecting that the first server is enabled to receive the performance data, the client receives performance data preference information from the first server, and wherein the performance data preference information comprises at least one of: under what circumstances the client should send

performance data, whether the client should send performance data related to communications with servers other than the first server, and a maximum performance data age threshold over which the client should not send performance data to the first server.